

Detecting non-binomial sex allocation when developmental mortality operates

RUNNING TITLE: Detecting non-binomiality

Richard D. Wilkinson^a, Apostolos Kapranas^{b,c}, Ian C.W. Hardy^b

^a*Corresponding author: School of Mathematics and Statistics, University of Sheffield, Sheffield, S3 7RH, UK. r.d.wilkinson@sheffield.ac.uk*

^b*School of Biosciences, University of Nottingham, Sutton Bonington Campus, Loughborough, LE12 5RD, UK*

^c*Laboratory for Fundamental and Applied Research in Chemical Ecology, University of Neuchtel, Neuchtel, 2000 Switzerland*

Abstract

Optimal sex allocation theory is one of the most intricately developed areas of evolutionary ecology. Under a range of conditions, particularly under population sub-division, selection favours sex being allocated to offspring non-randomly, generating non-binomial variances of offspring group sex ratios. Detecting non-binomial sex allocation is complicated by stochastic developmental mortality, as offspring sex can often only be identified on maturity with the sex of non-maturing offspring remaining unknown. We show that current approaches for detecting non-binomiality have limited ability to detect non-binomial sex allocation when developmental mortality has occurred. We present a new procedure using an explicit model of sex allocation and mortality and develop a Bayesian model selection approach (available as an R package). We use the double and multiplicative binomial distributions to model over- and under-dispersed sex allocation and show how to calculate Bayes factors for comparing these alternative models to the null hypothesis of binomial sex allocation. The ability to detect non-binomial sex allocation is greatly increased, particularly in cases where mortality is common. The use of Bayesian methods allows for the quantification of the evidence in favour of each hypothesis, and our modelling approach provides an improved descriptive capability over existing approaches. We use a simulation study to demonstrate substantial improvements in power for detecting non-binomial sex allocation in

situations where current methods fail, and we illustrate the approach in real scenarios using empirically obtained datasets on the sexual composition of groups of gregarious parasitoid wasps.

Key words: Sex ratio; under-dispersion; Bayes factor; Markov chain Monte Carlo

1. Introduction

The null model of sex allocation theory is the Düring-Fisher theory of equal investment (West, 2009). When populations are both large and have unbiased sex ratios, selection for variance in the sexual composition of offspring groups is predicted to be absent (Kolman, 1960). Under these conditions mothers will not be selectively penalized if they randomly allocate sex to offspring, with fixed probability of $p = 0.5$ that the offspring is male, independently of the sex of previous offspring. Thus, the number of males in each offspring group would have binomial variance, i.e., $np(1 - p)$, where n is the number of offspring. In smaller populations and under sex ratio bias ($p \neq 0.5$), stabilizing selection for low sex ratio variance is predicted, i.e., variance less than $np(1 - p)$ (Verner, 1965; West, 2009). Selection on sex ratio variance is likely to be strong when populations are structured into small reproductive subgroups within which offspring mate with each other on maturity and prior to the dispersal of the daughters (local mate competition; Hamilton, 1967); here, selection favours the evolution of low sex ratio variance, especially when one or a very few mothers contribute offspring to the locally mating group (Green et al., 1982; Hardy, 1992; Nagelkerke and Hardy, 1994; Nagelkerke, 1996; West and Herre, 1998). This is because low variance maximizes the production of mated daughters, a close correlate of maternal fitness. If one male is sufficient to mate successfully with all females within a group and all offspring in the group are progeny of one mother, then the optimal sexual composition is one male and the remainder of the group being females (Green et al., 1982). Similar arguments predict low variance under local resource competition (a generalization of local mate competition) and its converse, local resource enhancement (Lambin, 1994). Variance in the number of males among groups lower than expected under binomial sex allocation is known as under-dispersion, and sex allocation is then termed *precise* (Green et al., 1982; Lambin, 1994; Nagelkerke, 1996).

Control of sex allocation can be detected in some organisms by direct observation of sexually differential aspects of individual offspring production, such as maternal movements during egg laying, or the placement of offspring, or by non-random production sequences (Cole, 1981; Hardy, 1992; Heinsohn et al., 1997; Krackow et al., 2002; Khidr et al., 2013; Ambrosini et al., 2014) but such evidence is not often available. Empiricists must more frequently rely on the statistical analysis of offspring group sex ratios to detect whether sex allocation is being controlled or whether it is, for instance, binomial, as might be the null-expectation under several chromosomal mechanisms of sex-determination (Avilés et al., 2000; Krackow et al., 2002; Ewen et al., 2003; Macdonald and Johnson, 2008; Postma et al., 2011). Furthermore, empirical evaluations of sex ratio variance can provide tests of explicit predictions of sex ratio theory (e.g., Lambin, 1994; Morgan and Cook, 1994; Hardy and Cook, 1995; Hardy et al., 1998; Nagelkerke and Sabelis, 1998; West and Herre, 1998; Kapranas et al., 2011; Khidr et al., 2013; Bowers et al., 2013).

One practical problem often faced by investigations of sex ratios and sex ratio variance is that information on the sexual compositions of offspring is available at maturity but not at the time of sex allocation, and it is not uncommon for some offspring to die before maturity, (e.g., Hardy et al., 1998; Dyrce et al., 2004; Ewen et al., 2004; Forsyth et al., 2004; Dietrich-Bischoff et al., 2006; Øigarden and Lifjeld, 2013). Provided it has a stochastic component, developmental mortality will act to increase the variance of observed sex ratios, making initially under-dispersed data appear closer to binomial. This effect is expected on logical grounds (Section 3) and has been shown empirically both within and across several species of organisms with group structured mating (Hardy et al. 1998; Kapranas et al. 2011; Khidr et al. 2013; see also Dyrce et al. 2004 and Dietrich-Bischoff et al. (2006)). Current statistical approaches to assessing sex ratio variance (Krackow et al., 2002) are, however, based on the implicit assumption that developmental mortality does not operate, and they consequently lack power to detect non-binomiality, unless mortality rates are low.

Our aim is to show that by introducing a model that represents the biological processes that generated the data (sex allocation followed by mortality), we can substantially improve our ability to detect underlying biological behaviours. We also demonstrate the advantage of using more descriptive statistical approaches such as estimating effect sizes (with measures

of confidence), rather than relying on null-hypothesis significance testing, where the small dataset sizes mean we often fail to clear an arbitrary significance hurdle (usually $\alpha = 0.05$) even when the data indicate phenomena of interest. We begin by evaluating the performance, under developmental mortality, of the statistical methods commonly used to detect non-binomial sex ratio variance. We find that the power of these methods is adversely affected by developmental mortality. We then develop an alternative approach that explicitly models the mortality process. This has much improved power for detecting non-binomial sex allocation, particularly when there is high mortality or datasets are small.

2. Terms and notation

We define some terms and notation before describing current approaches and their limitations, and then introduce our new approach for detecting non-binomial sex allocation. A summary of the notation is provided in Table 1. The methods developed are general, but are likely to most readily be applied to egg-laying organisms such as birds, parasitoid wasps, fig wasps and phytoseiid mites (Hardy, 1992; Nagelkerke and Sabelis, 1998; West and Herre, 1998; West, 2009; Bowers et al., 2013), and this is reflected in the terminology we adopt (for a mammalian example see Macdonald and Johnson, 2008). Assume that we have a dataset containing data on C different clutches of eggs, all of which were laid in comparable environmental conditions. Offspring group size is called *clutch size* at the time of production (egg-laying) and *brood size* at the time of offspring maturity: brood size is less than clutch size when developmental mortality occurs.

A *primary* dataset consists of counts of the number of eggs and their sex for each clutch. Let N_i denote the number of eggs laid in the i^{th} clutch, and M_i be the number of those N_i eggs that are male. A primary dataset is the collection $\{(N_i, M_i)\}_{i=1}^C$. However, for most empirical investigations M_i is not observed, as the sex of an offspring cannot be easily determined from the eggs: it is usual to wait until the eggs hatch and develop to the point at which offspring sex can be discriminated (e.g., Dietrich-Bischoff et al., 2006; Khidr et al., 2013). It is also usual that a proportion of the eggs fail to mature, due to some form of developmental mortality, and consequently their sex cannot be recorded.

A *secondary* dataset consists of counts of n_i , the number of offspring that reach maturity

(brood size) and m_i , the number of those offspring that are male, with the complete secondary dataset denoted $\{(n_i, m_i)\}_{i=1}^C$. Although a small number of experiments have been conducted where primary datasets are obtained, either directly from genetic characteristics of eggs (Dijkstra, 1986; Hardy et al., 1998; Nagelkerke and Sabelis, 1998; Khidr et al., 2013) or through selective statistical procedures (Dyrce et al., 2004; Kapranas et al., 2011), the vast majority of analyses have been conducted using secondary datasets (e.g., Hardy, 1992; West and Herre, 1998; Nagelkerke and Sabelis, 1998; Mackauer and Völkl, 2002; Dietrich-Bischoff et al., 2006; Kapranas et al., 2008).

Our null hypothesis about sex allocation, H_0 , is that there is a sex ratio p (the proportion of offspring that are male), and that each egg is male with probability p independently of all other eggs in the clutch, i.e., that the distribution of sex ratios is binomial

$$M_i \sim \text{Bin}(N_i, p). \quad (1)$$

The alternative hypothesis, H_1 , is that the number of males is non-binomially distributed, that is, either over- or under-dispersed when compared to the binomial distribution. Note that these are hypotheses about primary sex ratios, not secondary sex ratios.

3. Current approaches for detecting non-binomial sex allocation

Several methods have been used for the statistical analysis of sex ratio variances (James, 1975; Green et al., 1982; Nagelkerke and Sabelis, 1998; West and Herre, 1998; Krackow et al., 2002). Whilst these methods can work well when applied to primary sex ratio data, this is not usually available, and so these methods are instead applied to secondary data, effectively treating them as if they were primary data. Not considering or ignoring that mortality has occurred thus violates the assumptions behind each approach; this results in a lack of statistical power, often leading to incorrect conclusions.

The first method for detecting departures from the binomial distribution, is a formal statistical test derived by E. Meelis (Nagelkerke and Sabelis, 1998), which we refer to as the *Meelis test* (Krackow et al., 2002). The test is a comparison of the estimated variance with the variance under the assumption of a binomial distribution, and is derived by calculating the distribution (under the null hypothesis) of $\sum m_i^2$ conditional on $\sum m_i$. A test statistic

Symbol	Definition
C	Number of clutches in the dataset
N	Number of eggs laid (primary)
M	Number of eggs laid that are male (primary)
n	Number of offspring that reach maturity (secondary)
m	Number of males that reach maturity (secondary)
D	The complete observed dataset, i.e., $D = \{(n_i, m_i)\}_{i=1}^C$
p	Sex ratio† (proportion of eggs that are male)
ψ	Dispersion parameter
λ	Average clutch size
d	Mortality rate
H_0, H_1	Null and alternative hypotheses
U	Test statistic for the Meelis' test
R	Descriptive ratio contrasting observed and expected variance
s^2	McCullagh's dispersion estimator
\mathcal{S}	Clutch sizes observed in the data, i.e., $\{k : n_j = k \text{ for some } j\}$
v_k	Number of clutches of size k , i.e., $\sum_{i=1}^C \mathbb{I}_{n_i=k}$
s_k^2	Empirical variance of the number of males in clutches of size k
B_{01}	Bayes factor for comparing H_0 with H_1

Table 1: Summary of notation used in this article. Letters in **bold** font indicate vector quantities, indices (e.g., n_i) indicate an instance of that variable, and hats (e.g., \hat{p}) indicate estimates. †Care needs to be taken with interpretation of p in the multiplicative binomial model as p is no longer the expected sex ratio when $\psi \neq 0$.

U (see supplementary material for details) is defined which can be shown to have a standard normal distribution under H_0 , provided C is sufficiently large. Large negative values of U indicate under-dispersion, and large positive values over-dispersion; typically, the test is applied by calculating the p-value $\mathbb{P}(|U| > |u_{obs}|)$, where \mathbb{P} denotes probability, with small values taken to indicate departure from the null hypothesis.

There are several difficulties with applying the Meelis test to the datasets used in empirical studies of sex-allocation. Firstly, the test assumes that the binomial random variables are observed directly, which is not the case when using secondary data (using m_i instead of M_i). Secondly, the test is derived for use on random variables from a binomial distribution with fixed size ($n_i = n$ for all i), whereas for real data, the values of n_i vary between broods, with datasets typically consisting of a range of brood sizes. It is common practice to collect all the broods of a certain size (e.g., all m_i such that $n_i = j$), then calculate the U -statistic, denoted U_j for those broods, before combining them using

$$U = \frac{\sum U_j}{\sqrt{|\mathcal{S}|}}$$

to give a single statistic U , where $\mathcal{S} = \{k : n_j = k \text{ for some } j\}$ is the collection of clutch sizes observed in the dataset. If each $U_j \sim N(0, 1)$, then $U \sim N(0, 1)$. However, the Meelis test was derived for large sample sizes. In practice, there may only be a small number of clutches with $n_i = j$, and so each U_j may not be well approximated by a standard normal distribution and hence, U may not have a $N(0, 1)$ distribution either.

James' test (James, 1975) is an alternative to the Meelis test that is often used for analysing datasets containing small clutches of unequal sizes. It involves calculation of a test statistic (Krackow et al., 2002, give details), which is known to be approximately normally distributed under the assumption of binomial sex ratios (no mortality). Large positive values indicate over-dispersion, and negative values under-dispersion. It is known to be less powerful for a single clutch size than the Meelis test (and suffers from the same difficulties as the Meelis test), but is included in our analysis for completeness.

The descriptive ratio R is also used:

$$R = \frac{\sum_{k \in \mathcal{S}} v_k s_k^2}{\sum_{k \in \mathcal{S}} v_k k \hat{p}_k (1 - \hat{p}_k)}$$

where s_k^2 is the empirical variance of the number of males in clutches of size k , i.e., $s_k^2 = \text{Var}(\{m_i : n_i = k\})$, and $v_k = \sum_{i=1}^C \mathbb{I}_{n_i=k}$ is the number of clutches which have size k . The denominator is the sum of the variances if assuming a binomial distribution, where \hat{p}_k is the estimated sex ratio for clutches of size k , i.e.,

$$\hat{p}_k = \frac{1}{kv_k} \sum_{i=1}^C m_i \mathbb{I}_{n_i=k}.$$

The rationale for using R , is that it is the observed variance of the number of males, divided by the variance that would occur if the number of males was binomially distributed (Krackow et al., 2002). We expect to observe $R \approx 1$ if the data are binomially distributed, with $R < 1$ for under-dispersed data. McCullagh and Nelder (1989) introduce a further estimator of dispersion, which is a sum of ratios rather than a ratio of sums

$$s^2 = \frac{1}{C-1} \sum_{i=1}^C \frac{(m_i - \hat{p}n_i)^2}{n_i \hat{p}(1 - \hat{p})} \quad \text{where} \quad \hat{p} = \frac{\sum m_i}{\sum n_i},$$

and should be interpreted in the same way as the R statistic.

The effect of mortality is to make the data appear less under-dispersed (more binomial), as mortality has the effect of increasing the variance of the number of males. To see this, imagine a species which has perfect precision, with each mother laying the same number of male and female eggs every time, so that the sex ratio variance is zero. Stochastic mortality would introduce an element of randomness to the sexual composition of the offspring groups, such that secondary datasets may even resemble binomial random variables under sufficiently high rates of mortality (see Section 5.3).

3.1. Evaluation of current approaches when developmental mortality occurs

To illustrate the limitations of current approaches, we simulate synthetic under-dispersed primary datasets, and then simulate the mortality process to produce synthetic secondary datasets. By applying the approaches described above, and repeating the process numerous times, we can examine their performance under varying levels of mortality.

We simulated sample experimental datasets as follows: for $i = 1, \dots, C$,

1. Simulate the clutch size from a Poisson distribution: $N_i \sim Po(\lambda)$, where λ is the average clutch size.

2. Simulate the number of males in the i^{th} clutch, M_i , from an under-dispersed multiplicative binomial distribution (Section 4).
3. Simulate the secondary dataset by assuming each of the N_i eggs has probability d of not reaching maturity, and count the number of females and males that survive.

We used a primary dataset on the parasitoid wasp *Goniozus legneri* (Khidr et al., 2013), a species known to produce a strongly under-dispersed primary sex ratio, to estimate parameter values for the synthetic data model, and used these estimates fixed throughout the simulation study ($\lambda = 10.0$, $p = 0.00278$, and $\psi = 0.445$, where p and ψ are parameters in the multiplicative binomial distribution, which is an under-dispersed distribution - see Section 4.1). We varied the size of the simulated experiment C , and the mortality rate d , and for each pair of values we simulated 10,000 synthetic datasets, and averaged the test statistics found across the replicates. This allows the effectiveness of all the procedures to be examined across a range of dataset sizes, C , and mortality rates d .

The performance of a hypothesis test can be measured by its power for a given significance level, where power is the probability of detecting non-binomial sex allocation when it occurs (i.e., $\text{power} = 1 - \mathbb{P}(\text{Type II error}) = \mathbb{P}(\text{reject } H_0 \mid H_1)$). Contour plots of the power of the Meelis and James tests (at significance level 0.05) as a function of the number of clutches in the dataset and the mortality rate show that the test lacks power if the number of clutches used is small or if the mortality rate is moderate-to-large (Fig. 1a,b). For example, for a dataset containing 50 clutches with a mortality rate of 10% there is only a 35% probability of correctly detecting under-dispersion. The power of the test is lower still if lesser degrees of under-dispersion are assumed as it becomes harder to detect (we used reasonably large under-dispersion of $\psi = 0.445$).

Fig. 1c shows the effect of mortality on R . The expected value of R increases towards 1 as the mortality rate increases, so that species with a high mortality rate will have R values consistent with binomial sex allocation, even if their primary sex ratios are under-dispersed. Fig. 1d shows the same information for McCullagh's s^2 . This can be seen to be less affected by mortality and so its use should be preferred to R . The number of clutches in the experiment has only a minor effect on the expected value of both statistics. However,




Figure 1: a) and b) show contour plots of the power of the two-sided Meelis and James tests; c) and d) are contour plots of the values of descriptive statistics R and McCullagh's s^2 , all as a function of C and d . The values were estimated using 10,000 randomly generated datasets, using parameter values estimated from data on *G. legneri* primary sex ratios.

it strongly affects the variance of the estimate (not shown), and for smaller experiments the observed values can vary widely, and so without appropriate confidence intervals for both statistics, they have little value.

There are two (related) reasons for the lack of power in these approaches. The first is that mortality increases the variance of the secondary values (n_i, m_i) compared to the primary values (N_i, M_i) making under-dispersion harder to detect. The second is that the tests do not take into account the fact that mortality has occurred, and consequently the additional variance is incorrectly interpreted as being consistent with binomial sex ratios.

4. A new test for detecting non-binomial sex allocation

By explicitly modelling mortality we develop a test that has improved statistical power as well as an increased descriptive capability. Our null hypothesis is a binomial model of sex allocation, which we compare to two different generalisations of the binomial distribution, the multiplicative binomial and the double binomial distributions, both of which can model over- and under-dispersion. Our model for the data then consists of a mortality model applied to the output of the sex allocation model. We use Bayesian model selection to determine which model is best supported by the data. The more intricate computational details are given in the supplementary material; here we focus on the broad outline of the approach.

4.1. A model of secondary data

We assume we have data on C different broods from comparable environmental conditions, so that they can be considered to be statistically exchangeable. Note that the unobserved primary counts N_i and M_i , and the corresponding secondary values after mortality has occurred, n_i and m_i , must satisfy the inequalities

$$N_i \geq n_i, \quad M_i \geq m_i \quad \text{and} \quad N_i - n_i \geq M_i - m_i. \quad (2)$$

We consider three models for the data, which differ only in the distribution of the sex allocation, i.e., the distribution of M_i given N_i . The first is the binomial model, with $M_i|N_i, p \sim \text{Bin}(N_i, p)$, which corresponds to the null hypothesis in Section 2. The second is

the multiplicative binomial distribution introduced by Altham (1978):

$$\mathbb{P}(M|N, p, \psi) = c(p, \psi) \binom{N}{M} p^M (1-p)^{N-M} e^{\psi M(N-M)}, \quad (3)$$

where $c(p, \psi)$ is an intractable normalising constant. The two parameters are a probability p , and a dispersion parameter ψ . The third, introduced by Efron (1986), is the double binomial model

$$\mathbb{P}(M|N, p, \psi) = c(p, \psi) \binom{N}{M} \frac{N^{N\psi} p^{M(\psi+1)} (1-p)^{(N-M)(\psi+1)}}{M^{M\psi} (N-M)^{(N-M)\psi}} \quad (4)$$

where $c(p, \psi)$ is again an intractable normalising constant. Note that when $\psi = 0$, both the multiplicative and double binomial distributions reduce to the binomial distribution. These models are the key part of our procedure, corresponding to the alternative hypothesis in Section 2, as they both model the three cases of interest:

- (i) binomial sex allocation when $\psi = 0$
- (ii) under-dispersed sex allocation when $\psi > 0$
- (iii) over-dispersed sex allocation when $\psi < 0$.

Unfortunately neither of these two distributions arises from a simple physical mechanism. Familiarity does allow an intuition to develop about the meaning of ψ , but our usage here does not require any interpretation beyond that given above, and that larger values of ψ indicate more under-dispersion than small values etc. Care also needs to be taken with interpretation of p , as the expected value of M is no longer Np for the multiplicative binomial distribution, except when $\psi = 0$, and so p can no longer be considered to be the sex ratio (the expected sex ratio, $\mathbb{E}\left(\frac{M}{N}\right)$, can be determined by Monte Carlo integration). We include both models as alternatives, as different datasets fit different models better, and this makes the detection of under-dispersion more likely.

We use the same model of mortality in each hypothesis and assume that each egg has probability d of dying before maturity, and thus of not being counted in the secondary dataset, independently of its sex and the other eggs in the clutch, i.e., we assume mortality is binomially distributed:

$$n_i|N_i, d \sim \text{Bin}(N_i, d). \quad (5)$$

The distribution of m_i can then be shown, by a label permuting argument, to be

$$\mathbb{P}(m|M, N, n, d) = \frac{\binom{M}{M-m} \binom{N-M}{N-n-M+m}}{\binom{N}{N-m}}. \quad (6)$$

We use two complimentary approaches for detecting departures from binomial sex allocation, the first based on estimation of effect size, and the second on hypothesis testing (Nakagawa and Cuthill, 2007). The simpler approach is to estimate the effect size, measured by the dispersion parameter ψ , by finding its posterior distribution $\pi(\psi|D)$. This parameter indicates whether sex allocation is binomial, over-, or under- dispersed, as well as how strong the effect is. Posterior credibility intervals for ψ can be used to assess the precision of the estimates and indicate informally whether the data are consistent with H_0 ($\psi = 0$). We describe methodology to do this below, the code is provided in the **precision** R package, and applications are described in Section 5.

While various authors recommend estimation over hypothesis testing (Robert, 2001; Gelman et al., 2003; Nakagawa and Cuthill, 2007), relying solely on estimation of ψ does not always provide the clarity required. For example, if the posterior distribution contains some support for $\psi = 0$, but the posterior mode is not close to 0, it can be difficult to judge whether or not data are under-dispersed using only the posterior distributions (Section 5.2). Instead, we wish to obtain the probability that sex allocation is under-dispersed, i.e., the posterior probability that each of the three models M_0 , M_1 and M_2 are true conditional upon the data: $\mathbb{P}(M_0|D)$, $\mathbb{P}(M_1|D)$, and $\mathbb{P}(M_2|D)$. These probabilities only make sense in a Bayesian setting, although note that p -values obtained from classical hypothesis tests, such as the Meelis test, are often incorrectly interpreted in this way (Goodman, 2008).

Bayesian model selection requires calculation of the Bayes factor (Jeffreys, 1939; Kass and Raftery, 1995), which is defined as the ratio of the evidence for two different hypotheses (or models)

$$B_{01} = \frac{\pi(D|H_1)}{\pi(D|H_0)}. \quad (7)$$

Values of B_{01} greater than 1 indicate evidence in favour of H_1 (over H_0) and values less than 1 indicate evidence for H_0 (over H_1). Jeffreys (1939) suggested interpretation of the strength of evidence in favour of a hypothesis according to the magnitude of the Bayes factor is shown

B_{01} range	$\mathbb{P}(H_1 D)$ range	Interpretation
1–3	0.5-0.75	Barely worth mentioning
3–10	0.75 - 0.91	Substantial
10–30	0.91-0.97	Strong
30–100	0.97- 0.99	Very strong
> 100	0.99-1	Decisive

Table 2: Jeffreys’ suggested interpretation of the Bayes factor for strength of evidence in favour of H_1 over H_0 . Values of $1/B_{01} = B_{10}$ give the strength of evidence for H_0 over H_1 . Also shown are the corresponding ranges of the posterior probability for H_1 given the data, in the case where we assign equal prior probability to both hypotheses.

in Table 2. The Bayes factor $B_{ij} = \mathbb{P}(D|H_j)/\mathbb{P}(D|H_i)$ can also be interpreted by noting that it is the ratio between the posterior and prior odds in favour of H_j over H_i

$$\frac{\mathbb{P}(H_j|D)}{\mathbb{P}(H_i|D)} = B_{ij} \frac{\pi_j}{\pi_i}$$

where π_j is the prior probability of H_j . Table 2 contains the posterior probabilities of H_1 being true for various Bayes factor ranges when we assume the hypotheses are equally likely *a priori*.

Bayes factors provide a powerful alternative to frequentist hypothesis tests, and have several advantages over classical methods. The first is that they provide a way to evaluate the evidence in favour of a hypothesis, in contrast to the classical approach which only rejects or accepts the null hypothesis for a particular error rate. This is particularly useful in datasets where the effect size or the sample size are small, or where the mortality rate is high, as we can quantify the strength of the evidence for under-dispersion in the data, even if there is not enough evidence to formally reject the null hypothesis. For instance, for analysis of data on *Goniozus thailandensis* (Section 5.2), the Meelis test finds $p > 0.05$ and thus concludes that there is no evidence to reject the null hypothesis, whereas the Bayesian approach reports that the posterior probability of the double binomial model being the true model is 0.79, with the probability of the binomial model being true only 0.14. When combined with the posterior distribution of ψ , which is concentrated on values greater than 0, this strongly suggests that this species produces under-dispersed sex ratios, a message that is lost if we only report the decision from the Meelis test.

4.2. Parameter estimation

We now describe how to find the posterior distribution of the parameters $\theta = (\psi, p, d, \lambda)$ given the data $D = \{(n_i, m_i)\}_{i=1}^C$, which we denote $\pi(\theta|D)$. This distribution represents our beliefs about the parameters after training the model to take the observed experimental data into account. The posterior distribution cannot be found analytically, and so we use Markov Chain Monte Carlo (MCMC) methods (e.g., Gilks et al., 1996) to obtain an approximation. We describe the case where only the n and m values, the number of eggs that reached maturity, have been recorded. The simpler situation where N_i is observed is a special case and follows immediately.

We introduce prior distributions for all unknowns. We assume the number of eggs laid in each clutch follows a Poisson distribution with mean λ

$$N_i \sim Po(\lambda) \text{ for } i = 1, \dots, C \quad (8)$$

and for the fixed parameters we assume that

$$\begin{aligned} p &\sim U[0, 1] & \psi &\sim N(0, \sigma^2) \\ \lambda &\sim \Gamma(a, b) & d &\sim \text{Beta}(a', b'). \end{aligned} \quad (9)$$

The distribution of p , λ and d are conjugate to the likelihood, allowing a Gibbs sampler to be used. Informative prior distributions are usually available for λ and d , as scientists often have information about mortality rates and average clutch sizes for the species of interest, although simulation suggests that the quantities of interest (the Bayes factor and the posterior of ψ), are robust to the choice of priors for λ and d . The key parameter is the dispersion parameter ψ , which we assign a zero mean normal distribution, so that under- and over-dispersion are equally likely *a priori*. We use an uninformative prior distribution for p , so that the posterior distribution is determined solely by the data.

To sample from the posterior distribution $\pi(\theta|D)$, we use a Metropolis-Hastings within Gibbs sampler (Metropolis et al., 1953; Geman and Geman, 1984). We introduce vectors of unobserved N_i and M_i values, denoted \mathbf{N} and \mathbf{M} , as auxiliary variables, and sample across the chain $\pi(\theta, \mathbf{N}, \mathbf{M}|D)$, which is a $(4 + 2C)$ dimensional Markov chain. The distribution of interest, $\pi(\psi|D)$, is then found by taking the marginal distribution of ψ . Details of the

MCMC algorithm used are provided in the supplementary material, and the algorithm is implemented in the accompanying **precision** R package for each of the three models.

4.3. Bayes factor estimation

To calculate the Bayes factors (Equation 7) we must first calculate the evidence for each model

$$\pi(\mathbf{n}, \mathbf{m}) = \int \pi(\mathbf{n}, \mathbf{m}|\theta)\pi(\theta)d\theta,$$

where \mathbf{n} and \mathbf{m} are the vectors of the observed n_i and m_i values, which is analytically intractable for the models considered. We use the approach described in Chib (1995) and Chib and Jeliazkov (2001) to estimate the evidence for each model, which relies upon the identity

$$\pi(\mathbf{n}, \mathbf{m}) = \frac{\pi(\mathbf{n}, \mathbf{m}|\theta^*)\pi(\theta^*)}{\pi(\theta^*|\mathbf{n}, \mathbf{m})}. \quad (10)$$

Calculation of both the numerator and denominator is challenging, but can be done with additional samples from an MCMC sampler. The derivation and details of the algorithm are technical, and are presented in the supplementary material. An implementation of these algorithms is available as the **precision** R package, available on github. The next section demonstrates the power of our approach.

5. Results

We illustrate our approach using data on four species of wasp: The strength of evidence for under-dispersion from secondary sex ratio data in these species varies from weak (*Colpoclypeus florus*) to overwhelming (*Metapycus luteolus*), and the mortality rate varies from low (*Goniozus legneri*) to high (*C. florus*). We also present the results from a simulation study which conclusively demonstrates the increased power of our approach.

The Bayesian approach requires prior distributions for all unknown parameters. Simulation studies have shown that the model and data are strongly informative about p and ψ , so that any information in the prior distribution is overwhelmed by the information in the data. In all our analyses we give p an uninformative prior distribution uniform on $[0, 1]$ and ψ a vague prior distribution for both the double and multiplicative binomial models:

$$p \sim U[0, 1], \quad \psi \sim N(0, 1). \quad (11)$$

The prior for ψ can be justified by examining the degree of under-dispersion for various levels of ψ . If $M \sim \text{DoubleBinom}(n = 10, p = 0.1, \psi)$, then $\mathbb{P}(M = 1) = 0.38$ if $\psi = 0$ (the binomial case), whereas for $\psi = 3$, $\mathbb{P}(M = 1) = 0.85$, indicating strong under-dispersion. The Bayes factors are robust to the choice of priors for λ, d and p (the parameters shared across models), but unsurprisingly, are sensitive to the prior for ψ . More diffuse priors for ψ tend to reduce the evidence for under-dispersion due to an Occam’s razor type effect, but for realistic priors for ψ , the conclusion does not usually change significantly (see supplementary material). Fortunately, the posterior distribution for ψ is robust to the choice of prior for ψ , and so this can also be used to indicate whether the data are under-dispersed.

The data typically contain only limited information about the parameters λ and d , but with the two posterior distributions strongly correlated, as large average clutch size and high mortality, or small average clutch size and lower mortality rate, leads to similar datasets. Prior information about λ and d is often available, which we can use to choose prior distributions for these two parameters on a species by species basis. Experimentation has shown that the Bayes factor and the posterior distribution of p and ψ (the primary parameter of interest) are robust to these choices.

5.1. *Goniozus legneri*: Large dataset, low mortality

We begin by considering data on *G. legneri*, a gregarious parasitoid wasp in which offspring groups are produced by single mothers and sex ratios are female biased due to local mate competition. Khidr et al. (2013) provide both a primary dataset, consisting of pre-mortality counts on 47 clutches obtained using DNA microsatellite markers to identify the sex of eggs, and a secondary dataset containing post-mortality counts of male and female adults in 113 clutches. Both the Meelis and James tests lead to rejection of the null hypothesis of binomial sex allocation (Table 3) with p -values of 0.0041 and 0.0027 respectively for the secondary data. Furthermore, we find $R = 0.572$, which when combined with the negative value of U in the two tests ($U = -2.38$ and $U = -1.98$ for Meelis and James respectively), lead us to conclude, in common with previous studies (Hardy et al., 1998; Khidr et al., 2013), that *G. legneri* has under-dispersed sex ratios.

Khidr et al. (2013) reported that the proportion of offspring that died before maturity was 7.6%, which agrees with previous *G. legneri* mortality estimates (5-12%, Hardy et al., 1998).

Species									
		<i>G. legneri</i>		<i>G. thailandensis</i>		<i>C. florus</i>		<i>M. luteolus</i>	
Procedure	Instance	Primary Value	Secondary Value	Secondary Value	Primary Value	Secondary Value	Primary Value	Secondary Value	
James	U	-1.98	-3.00	-2.01	-0.89	2.7	-6.7	-7.8	
	p	0.047	0.0027	0.045	0.37	0.0068	1.5×10^{-11}	7.0×10^{-15}	
Meelis	U	-2.38	-2.87	-0.73	-3.24	-0.97	-7.9	-7.4	
	p	0.017	0.0041	0.46	0.0012	0.33	2.6×10^{-15}	1.3×10^{-13}	
<i>R</i>		0.44	0.57	0.68	0.13	0.75	0.093	0.44	
<i>s</i> ²		0.57	0.61	0.74	0.51	1.18	0.20	0.58	
BF	double:binomial	45.1	213.6	5.65	3830	0.27	1.1×10^{29}	9.8×10^{23}	
	multiplicative:binomial	9430	31.3	0.54	0.36	0.36	7.0×10^5	2.0×10^6	
	double:multiplicative	0.0048	6.8	10.5	10600	0.74	1.6×10^{23}	5.0×10^{17}	
Posterior probability	binomial	0.00010	0.004	0.14	0.00026	0.61	0.000	0.000	
	multiplicative	0.995	0.127	0.074	0.000094	0.22	0.000	0.000	
	double	0.0048	0.869	0.74	0.9996	0.16	1.000	1.000	

Table 3: Analysis of four wasp datasets: *G. legneri* primary ($C = 47$) and secondary ($C = 113$) datasets (Khidr et al., 2013); *G. thailandensis* secondary dataset ($C = 60$) (Witethom and Gordh, 1994); *C. florus* primary ($C = 55$) and secondary datasets ($C = 53$) (Dijkstra, 1986; Hardy et al., 1998); *M. luteolus* primary ($C = 127$) and secondary ($C = 371$) datasets (Kapranas et al., 2011). All values estimated using 10^6 MCMC iterations.

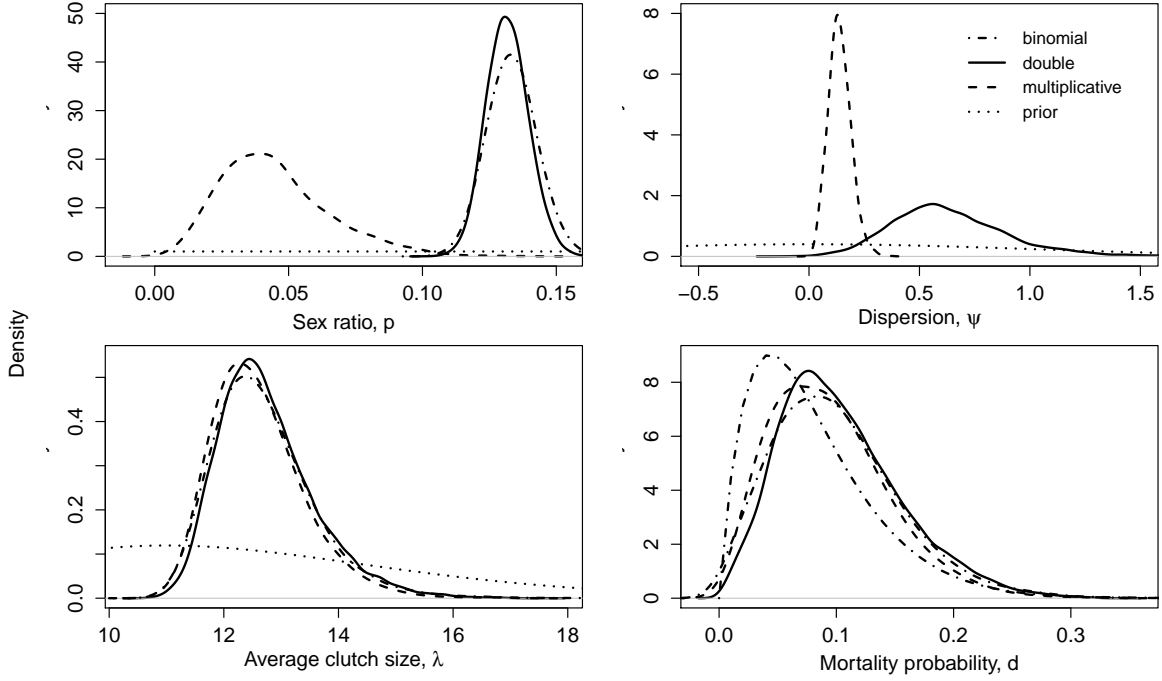


Figure 2: Posterior distributions from the analysis of *G. legneri* secondary data (Khidr et al., 2013), obtained using 5×10^5 MCMC iterations. For each parameter, the prior and posterior distribution are shown for the three alternative models of sex allocation. Note that the binomial model does not have a dispersion parameter (ψ) and that the interpretation of p and ψ is different in each model.

We incorporate this information into the analysis through the use of prior distributions

$$d \sim \text{Beta}(2, 23) \quad \lambda \sim \text{Gamma}(12, 1).$$

The prior mean for d is thus $2/(23 + 2) = 8\%$, with values in the range 0-20% all supported *a priori* (Figure 2). The prior for λ was based on an observed secondary clutch size of 11, and the mortality rate of 7.6%, suggesting a prior mean for λ of approximately 12. The $\text{Gamma}(12, 1)$ prior distribution has a prior mean of $12/1$, and supports prior λ values in a range between 11 and 14 (Figure 2).

Figure 2 shows the posterior distributions of the four parameters for the secondary dataset. Interest lies primarily in the dispersion parameter ψ , with $\psi > 0$ indicating under-dispersion and $\psi < 0$ over-dispersion. We cannot estimate ψ precisely as there is a finite quantity of data, but the posterior distributions show the range of ψ values we believe could feasibly have led to the observed data. The posterior distribution for ψ for both the double

and multiplicative models, suggests that only positive values of ψ are consistent with the data. Equi-tailed 95% credibility intervals for ψ are $[0.047, 0.248]$ for the multiplicative binomial model, and $[0.196, 1.28]$ for the double binomial model, neither of which overlap with 0, leading us to conclude that *G. legneri* has under-dispersed sex allocation.

The Bayes factor (BF) estimates for *G. legneri* are reported in Table 3. We find that the double binomial model is best supported, with a BF of 213.6 in favour of the double binomial over the binomial model, which Jeffreys' scale interprets as decisive evidence. There is also very strong evidence in favour of the multiplicative model over the binomial ($BF = 31.3$), and substantial evidence to suggest the double binomial is better supported than the multiplicative binomial model ($BF = 6.8$). If we are prepared to assign all three models equal prior probability, then the posterior probability that the binomial model is the true model is 0.004, compared to 0.869 for the double binomial model, and 0.127 for the multiplicative binomial model.

For this dataset, the signal from the data is strong ($C = 113$ is a reasonably large sample size), and consequently all the procedures give unambiguous conclusions. However, it is informative to note the difference between the two approaches: the Meelis test strongly rejects H_0 , but does not indicate the size of the effect (the R value does indicate the size of the effect, but is unreliable without a measure of uncertainty). The p -value does not give the probability that H_0 is true and should not be interpreted as such. Meanwhile, the Bayesian procedure estimates the probability that H_0 is true, and the posterior distribution for ψ gives the effect size after having accounted for mortality, along with a measure of the uncertainty in the estimate of ψ . For *G. legneri*, Khidr et al. (2013) also provide a primary dataset which we can analyse without modelling mortality (Table 3). The conclusion is the same as for the secondary data, again with strong evidence of under-dispersion. One difference between the primary and secondary analyses is that for the primary data, the multiplicative binomial model is preferred, whereas for the secondary data, the double binomial model is preferred. We believe this is due to differences between the shape of the two distributions. Figure 3 shows the posterior predictive distribution for the number of male eggs laid (in a clutch of 10 eggs) for the six different scenarios (three models on both the primary and secondary data). We can see that for a given sex allocation model, the posterior predictions for the primary




Figure 3: Posterior predictive distributions of the number of male eggs (pre-mortality) in a clutch of 10 eggs for the six different scenarios, namely the primary and secondary data for the three different models of sex allocation. The multiplicative binomial distribution gives the best fit to the primary data, and the double binomial distribution best fits the secondary data.

and secondary data are similar, and that the double and multiplicative distributions both give more concentrated (more precise) predictions than the binomial model. We can also see the difference between the shape of the double and multiplicative distributions, with the multiplicative distribution predicting more clutches with no males than the double binomial. The switch between preferred model for the secondary and primary datasets does not change our conclusion that there is strong evidence of under-dispersion.

Finally, note that the data and model are strongly informative about p and ψ , with the posterior and prior values being markedly different, whereas the posterior value for λ and d are close to the prior distribution. Experimentation (see the supplementary material) has shown that the posterior distributions of λ and d are sensitive to their prior distribution, but that the posterior of p and ψ are not sensitive to these choices.

5.2. *Goniozus thailandensis*: small dataset, medium mortality

Now we consider a dataset on the parasitoid species *Goniozus thailandensis* collected by Witethom and Gordh (1994). This species has a broadly similar biology to *G. legneri* and has previously been analysed for sex ratio variance by Hardy et al. (1998). The developmental

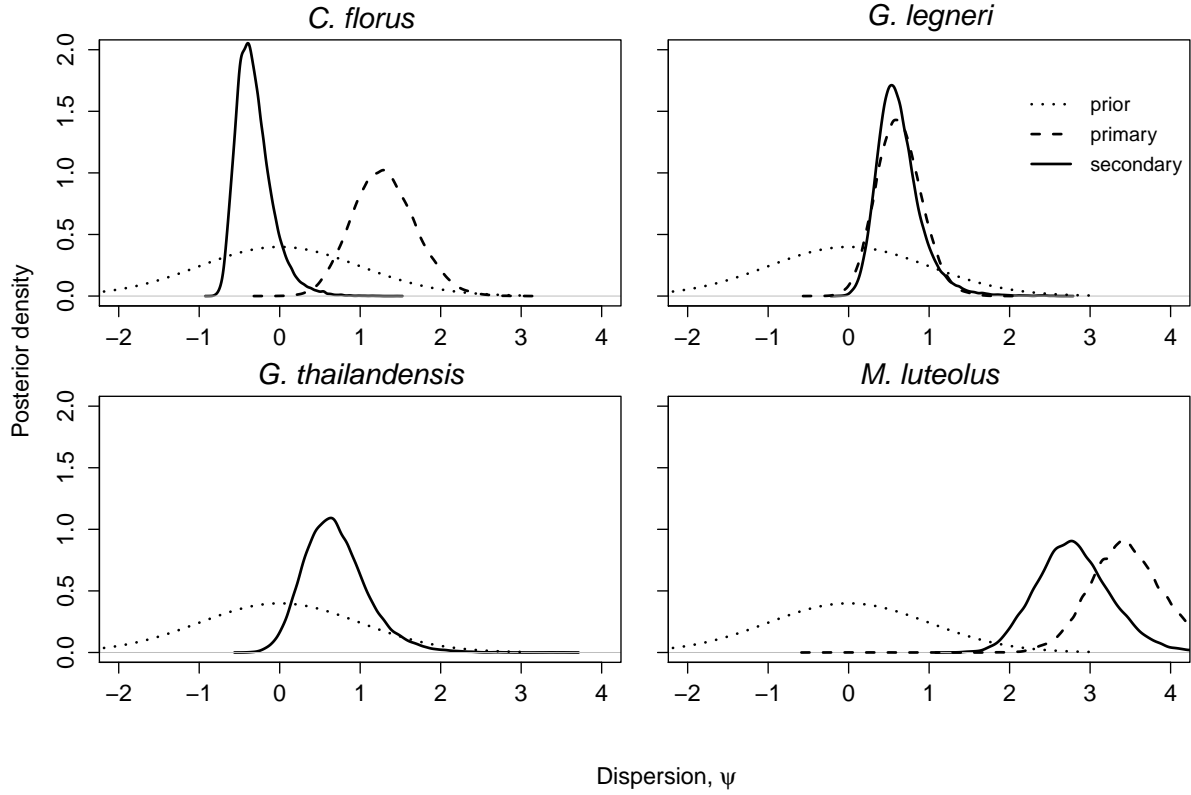


Figure 4: The marginal posterior distribution for ψ for the double binomial model for all four species. The results were obtained using 5×10^5 MCMC iterations. The posterior distributions for p , d and λ are not shown.

mortality rate, 22%, is higher than for *G. legneri* and the dataset is small, thus presenting a more challenging, and possibly more typical, case for analysis. Classical analysis of these data was inconclusive: the Meelis test gave $U = -0.73$ with a p -value of 0.23 and $R = 0.68$, which suggests under-dispersion, but with insufficient evidence to reject H_0 at the 5% significance level. In Section 3 we demonstrated that the Meelis test will lack power on this dataset, as there are only $C = 60$ observations and the probability of developmental mortality is moderate. This leaves us uncertain as to whether this result is due to the limited sample size, the relatively high mortality rate or to sex allocation actually being binomially distributed. The Meelis test only informs us that we cannot reject the null hypothesis due to insufficient evidence; it does not allow us to say that the species has binomially distributed sex allocation.

Carrying out the Bayesian analysis, using the prior distributions

$$d \sim \text{Beta}(5, 20) \quad \lambda \sim \text{Gamma}(9, 1),$$

(consistent with the observed average clutch size and the mortality rate of 22%) we find the posterior distribution for ψ shown in the bottom left panel of Figure 4 and the Bayes factors given in Table 3. The Bayes factors suggest that there is substantial evidence in favour of the double binomial model over the other two models, and the posterior for ψ shows that under-dispersion is the best explanation of the data (the equi-tailed 95% credibility interval for ψ is $[0.0415, 1.61]$). The posterior distribution does contain a small amount of support for a zero or negative value of ψ (binomiality, or over-dispersion), showing that while this can not conclusively be ruled out, it is unlikely. Assuming equal prior probability for each model, there is a posterior probability of 0.79 that the double binomial model is the true model, and 0.14 that the binomial model (H_0) is true. While this is not conclusive evidence, it has allowed us to state that the data suggest under-dispersion over binomial sex allocation with posterior odds of more than 5 to 1. The posterior for ψ allows us to see the range of possible under-dispersion strengths that are consistent with the data. In comparison, the classical approach only allows us to conclude that there is insignificant evidence to reject H_0 .

5.3. *Colpoclypeus florus*: medium dataset, high mortality

Primary and secondary data on *Colpoclypeus florus* are available from a study by Dijkstra (1986) analysed by Hardy et al. (1998). *C. florus* is a gregarious parasitoid with female biased sex ratios and is the only known member of its genus. The mortality rate was reported to be 57%, which when combined with the average clutch size of 7.4 motivated the prior distributions

$$d \sim \text{Beta}(11, 10) \quad \lambda \sim \text{Gamma}(16, 1).$$

The results of the analysis of this data are shown in Table 3. These illustrate the tendency of mortality to make data appear less under-dispersed, possibly even over-dispersed. The primary data clearly show that the species has under-dispersed sex allocation, with the Meelis test and Bayes factors agreeing that there is very strong evidence in favour of under-dispersion. Whereas for the secondary data, the Meelis test fails to reject the null hypothesis, and the Bayes factors suggest that the binomial model is the best supported (posterior probability of 0.61, compared to $0.16+0.22=0.38$ for the two non-binomial models). The 95% credibility interval for ψ is $[-0.063, 0.019]$ for the multiplicative model, and $[-0.65, 0.24]$ for

the double binomial model, both of which contain 0, showing that the data could be either under- or over-dispersed. The marginal posterior for ψ in Figure 4, shows how the primary data strongly suggest under-dispersion, but that the secondary data (after mortality) suggest over-dispersion, although there is still some support for under-dispersion. While the Meelis test can only lead us to conclude that there is no evidence to reject the hypothesis of binomial sex allocation, the Bayesian test can quantify that evidence and give a posterior probability that indicates that the hypothesis of binomial sex ratios is approximately twice as likely as the hypothesis of non-binomial sex allocation.

5.4. *Metaphycus luteolus*: large dataset, high mortality

A large secondary dataset on *M. luteolus* was presented in Kapranas et al. (2011). This species is a facultatively gregarious parasitoid which lays eggs inside hosts. Developing offspring may compete within the host, be attacked by the host immune responses, or die of other causes, and the overall mortality rate is approximately 40%. The secondary sex ratio is female biased. Using prior distributions

$$d \sim \text{Beta}(6, 10) \quad \lambda \sim \text{Gamma}(4, 1)$$

we obtained the results presented in Table 3 and Figure 4. Due to the large sample sizes, and the effect size, all procedures give overwhelming evidence that the data are under-dispersed. By selecting only those clutches that did not experience any mortality, we can obtain an approximation of a primary dataset (this approach is discussed in Khidr et al., 2013). Analysis of this dataset again demonstrates the tendency of mortality to make data appear less under-dispersed.

5.5. *Simulation study*

We now show that by modelling mortality, we have increased our ability to detect under-dispersion. We analyse the performance of the Meelis test and the Bayes factor approach, using a simulation study in which we apply both procedures to synthetic datasets. The computational expense of the Bayesian approach (typically it takes 2-5 hours of computer time to analyse a single dataset), limited the study to 100 synthetic datasets, but this is sufficient to conclusively demonstrate an improved ability to find evidence against H_0 , i.e., statistical power.




Figure 5: Results of the simulation study. Each point represents the p -value and Bayes factor (BF) for a simulated dataset. The shaded regions indicate p -value or Bayes factor ranges for which we would conclude there was either no, or weak evidence against H_0 . The horizontal regions (with ‘forwardslash’ shading) indicates the Bayes factor is either less than 10 (threshold for strong evidence against H_0), or less than 3 (threshold for substantial evidence). The vertical regions (‘backslash’ shading) indicate p -values of less than 0.01 or 0.05. Note that the x-axis is reversed.

The synthetic datasets were simulated to each contain 50 clutches using a mortality rate of 30%, moderate values of C and d . The model defined by Equations (3), (5) and (8), with $\lambda = 10$, $p = 0.1$, and $\psi = 0.3$, was used to simulate the datasets, giving a moderate level of under-dispersion comparable to *G. legneri*.

The results of the simulation study are summarised in Figure 5 and Table 4. For each dataset we have plotted the logarithm of the estimated Bayes factor between the multiplicative and binomial models, against the logarithm of the p -value from the Meelis test. The shading shows regions in which one or both of the procedures failed to detect strong evidence of under-dispersion, either because the p -value is greater than 0.05 (or 0.01), and/or because the Bayes factor is less than 3 (or 10). Table 4 summarises each procedure by the percentage of datasets which led to Bayes factors or p -values in a specified range.

Strength of evidence:	insubstantial	substantial	strong	very strong	decisive
Meelis p -value range:	> 0.1	$0.05 - 0.1$	$0.01 - 0.05$	$0.001 - 0.01$	< 0.001
% in range:	33	20	39	7	1
BF range:	$0 - 3$	$3 - 10$	$10 - 30$	$30 - 100$	> 100
% in range:	5	13	15	15	52

Table 4: Simulation study results: 100 synthetic datasets, all with moderate levels of under-dispersion ($\psi = 0.3$ in the multiplicative model) and mortality (30%), were analysed and grouped into categories indicating various levels of strength of evidence against H_0 . The Bayesian approach can be seen to substantially outperform the Meelis test.

These results clearly demonstrate the improved power of the Bayesian procedure. For example, in more than half of the simulated datasets, the Meelis test returned a p -value greater than 0.05, which would indicate that there was insufficient evidence to reject the null hypothesis of binomial sex allocation. In contrast, 95% of the datasets provided at least substantial evidence against binomial sex ratios according to the Bayesian approach, and over half (52%) of the datasets provided decisive evidence ($BF > 100$). Furthermore, Figure 5 illustrates that every time the Bayesian test failed to detect under-dispersion, the Meelis test also failed, whereas there were 36 datasets where the Bayesian test indicated strong evidence ($BF > 10$) against H_0 , but where the Meelis test failed (at the 5% level).

In order to confirm that this increased power is not due to a corresponding increase in the type I error rate (i.e., falsely rejecting H_0), a second simulation study was performed analysing synthetic datasets generated from the binomial model. For 200 simulated datasets, the Meelis test rejected H_0 (at $\alpha = 0.05$) in 3% of cases (i.e., it had approximately the assumed error rate). The Bayes factor gave $\mathbb{P}(H_0|D) \leq 0.05$ (i.e., strong evidence against H_0) in 6% of cases, showing that the increased power of the Bayesian approach is not due to an inflated type I error. The posterior distributions for ψ (available in the supplementary information), ruled out $\psi = 0$ in only one of the 200 simulated datasets.

6. Conclusions

We have shown that the current approaches used to detect under- or over-dispersion in sex allocation lack power when the sample size is small or the mortality rate is moderate to

large. Both are common situations in empirical studies. For example, the Meelis test will usually fail to reject the null hypothesis under these conditions even when sex allocation is strongly non-binomial. We have introduced a new approach to detecting under- or over-dispersion that has much greater power for detecting departures from binomial allocation. The approach gains its power by explicitly modelling mortality, so that the test takes into account that the patterns in the data have occurred through a combination of sex allocation and mortality. The method can be extended further to include non-binomial distributions of mortality (e.g., Hardy et al., 1998; Kapranas et al., 2011). Furthermore, using a Bayesian approach to model selection and parameter estimation increases our descriptive ability: the posterior distribution of the dispersion parameter ψ allows both the size of the effect and the range of possible effects that are consistent with the data to be identified. Using Bayes factors allows us to give the posterior probability that the data derive from a species that has binomially distributed sex allocation, as opposed to p -values, which although commonly interpreted as probabilities, should not be (Goodman, 2008). In situations where the evidence is conclusively in favour of one hypothesis, our test generates the same conclusion as current approaches (but with improved descriptive ability). However, when the evidence is weaker, the additional information provided by the Bayesian approach can allow us to make useful inferences, even if these cannot be conclusive.

7. Coda

The software implementing this approach has been written in R (R Development Core Team, 2008) and is freely available (<https://github.com/rich-d-wilkinson/precision>) as the **precision** R package on github. Details of how to use and install the package are given in the package vignette and in the supplementary material. There are many possible extensions to this approach, primarily through changes and improvements to the model. For example, the binomial mortality model is relatively simple and other more complex models (such as over-dispersion) are possible. These extensions are straightforward to make within the Bayesian testing framework.

The data used in this paper are all available within the **precision** R package (see the package vignette). These datasets, as well as additional data on the sexual compositions

of offspring groups, are available from several previous publications. Secondary sex ratio datasets can be found in Morgan and Cook (1994); Hardy and Cook (1995); Nagelkerke and Sabelis (1998); Mackauer and Völkl (2002); Kapranas et al. (2008, 2009) and Khidr et al. (2013). Primary sex ratios are more difficult to evaluate, but datasets are available in Dijkstra (1986); Avilés et al. (2000), and Khidr et al. (2013).

Acknowledgements: We thank S. K. Khidr, B. Witethom and L. J. Dijkstra for help with data. We thank Andrew Wood for useful advice on the statistical approach taken, and R. F. Green and three anonymous referees for constructive suggestions. Apostolos Kapranas was funded by a Marie Curie Fellowship (FP7-PEOPLE-IEF-273431).

References

- Altham, P. M. E., 1978. Two generalizations of the binomial distribution. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 27, 162–167.
- Ambrosini, R., Rubolini, D., Saino, N., 2014. Analysis of sex sequences by means of generalized linear mixed models. *Behavioral Ecology and Sociobiology* 68, 1367–1377.
- Avilés, L., McCormack, J., Cutter, A., Bukowski, T., 2000. Precise, highly female-biased sex ratios in a social spider. *Proceedings of the Royal Society of London B: Biological Sciences* 267, 1445–1449.
- Bowers, E. K., Munclinger, P., Bureš, S., Kučerová, L., Nádvorník, P., Krist, M., 2013. Cross-fostering eggs reveals that female collared flycatchers adjust clutch sex ratios according to parental ability to invest in offspring. *Molecular Ecology* 22, 215–228.
- Chib, S., 1995. Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association* 90, 1313–1321.
- Chib, S., Jeliazkov, I., 2001. Marginal likelihood from the metropolis-hastings output. *Journal of the American Statistical Association* 96, 270–281.
- Cole, L. R., 1981. A visible sign of a fertilization action during oviposition by an ichneumonid wasp, *itoplectis maculator*. *Animal Behaviour* 29, 299–300.

- Dietrich-Bischoff, V., Schmoll, T., Winkel, W., Krackow, S., Lubjuhn, T., 2006. Extra-pair paternity, offspring mortality and offspring sex ratio in the socially monogamous coal tit (*Parus ater*). *Behavioral Ecology and Sociobiology* 60, 563–571.
- Dijkstra, L. J., 1986. Optimal selection and exploitation of hosts in the parasitic wasp *Colpoclypeus florus* (Hym., Eulophidae). *Netherlands Journal of Zoology* 36, 177–301.
- Dyrce, A., Sauer-Gürth, H., Tkadlec, E., Wink, M., 2004. Offspring sex ratio variation in relation to brood size and mortality in a promiscuous species: the Aquatic Warbler *Acrocephalus paludicola*. *Ibis* 146, 269–280.
- Efron, B., 1986. Double exponential families and their use in generalized linear regression. *Journal of the American Statistical Association* 81, 709–721.
- Ewen, J. G., Cassey, P., King, R. A., Brittingham, M., 2003. Assessment of the randomization test for binomial sex-ratio distributions in birds. *The Auk* 120, 62–68.
- Ewen, J. G., Cassey, P., Møller, A. P., 2004. Facultative primary sex ratio variation: a lack of evidence in birds? *Proceedings of the Royal Society of London, Series B: Biological Sciences* 271, 1277–1282.
- Forsyth, D. M., Tustin, K. G., Gaillard, J.-M., Loison, A., 2004. Fetal sex ratio variation in the highly polygynous Himalayan tahr: evidence for differential male mortality. *Behavioral Ecology* 15, 572–578.
- Gelman, A., Carlin, J. B., Stern, H. S., Rubin, D. B., 2003. *Bayesian data analysis*. CRC press.
- Geman, S., Geman, D., 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6, 721–741.
- Gilks, W., Richardson, S., Spiegelhalter, D., 1996. *Markov Chain Monte Carlo in Practice*. Interdisciplinary Statistics. Chapman & Hall.

- Goodman, S., 2008. A dirty dozen: Twelve p-value misconceptions. *Seminars in Hematology* 45, 135–140.
- Green, R. F., Gordh, G., Hawkins, B. A., 1982. Precise sex ratios in highly inbred parasitic wasps. *American Naturalist* 120, 653–665.
- Hamilton, W. D., 1967. Extraordinary sex ratios. *Science* 156, 477–488.
- Hardy, I. C., Cook, J. M., 1995. Brood sex ratio variance, developmental mortality and virginity in a gregarious parasitoid wasp. *Oecologia* 103, 162–169.
- Hardy, I. C. W., 1992. Non-binomial sex allocation and brood sex ratio variances in the parasitoid Hymenoptera. *Oikos* 65, 143–158.
- Hardy, I. C. W., Dijkstra, L. J., Gillis, J. E. M., Luft, P. A., 1998. Patterns of sex ratio, virginity and developmental mortality in gregarious parasitoids. *Biological Journal of the Linnean Society* 64, 239–270.
- Heinsohn, R., Legge, S., Barry, S., 1997. Extreme bias in sex allocation in *Eclectus* parrots. *Proceedings of the Royal Society of London. Series B: Biological Sciences* 264, 1325–1329.
- James, W. H., 1975. Sex ratio and the sex composition of the existing sibs. *Annals of Human Genetics* 38, 371–378.
- Jeffreys, H., 1939. *Theory of probability*. International series of monographs on physics. The Clarendon press.
- Kapranas, A., Hardy, I. C. W., Morse, J. G., Luck, R. F., 2011. Parasitoid developmental mortality in the field: patterns, causes and consequences for sex ratio and virginity. *Journal of Animal Ecology* 80, 192–203.
- Kapranas, A., Pacheco, P., Forster, L., Morse, J. G., Luck, R. F., 2008. Precise sex allocation manifested by several encyrtid parasitoids of brown soft scale *Coccus hesperidum* L. (Hemiptera: Coccidae). *Behavioral Ecology and Sociobiology* 62, 901–912.

- Kapranas, A., Wajnberg, E., Luck, R. F., 2009. Sequences of sex allocation and mortality in clutches of *metaphycus* parasitoids of soft scale insects and the prevalence of all-female broods. *Ecological Entomology* 34, 652–662.
- Kass, R. E., Raftery, A. E., 1995. Bayes Factors. *Journal of the American Statistical Association* 90, 773–795.
- Khidr, S. K., Mayes, S., Hardy, I. C. W., 2013. Primary and secondary sex ratios in a gregarious parasitoid with local mate competition. *Behavioral Ecology* 24, 435–443.
- Kolman, W. A., 1960. The mechanism of natural selection for the sex ratio. *American Naturalist* 94, 373–377.
- Krackow, S., Meelis, E., Hardy, I. C. W., 2002. Analysis of sex ratio variances and sequences of sex allocation. In: Hardy, I. C. W. (Ed.), *Sex Ratios: Concepts and Research Methods*. Cambridge: Cambridge University Press, pp. 112–131.
- Lambin, X., 1994. Sex ratio variation in relation to female philopatry in Townsend’s voles. *Journal of Animal Ecology* 63, 945–953.
- Macdonald, D. W., Johnson, P. J., 2008. Sex ratio variation and mixed pairs in roe deer: evidence for control of sex allocation? *Oecologia* 158, 361–370.
- Mackauer, M., Völkl, W., 2002. Brood-size and sex-ratio variation in field populations of three species of solitary aphid parasitoids (Hymenoptera: Braconidae, Aphidiinae). *Oecologia* 131, 296–305.
- McCullagh, P., Nelder, J. A., 1989. *Generalized linear models*. Vol. 37. CRC press.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., Teller, E., 1953. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics* 21, 1087–1092.
- Morgan, D. J., Cook, J. M., 1994. Extremely precise sex ratios in small clutches of a bethylid wasp. *Oikos*, 423–430.

- Nagelkerke, C. J., 1996. Discrete clutch sizes, local mate competition, and the evolution of precise sex allocation. *Theoretical population biology* 49, 314–343.
- Nagelkerke, C. J., Hardy, I. C. W., 1994. The influence of developmental mortality on optimal sex allocation under local mate competition. *Behavioral Ecology* 5, 401–411.
- Nagelkerke, C. J., Sabelis, M. W., 1998. Precise control of sex allocation in pseudo-arrhenotokous phytoseiid mites. *Journal of Evolutionary Biology* 11, 649–684.
- Nakagawa, S., Cuthill, I. C., 2007. Effect size, confidence interval and statistical significance: a practical guide for biologists. *Biological Reviews* 82, 591–605.
- Øigarden, T., Lifjeld, J. T., 2013. Primary sex ratios vary with clutch size in the size-dimorphic White-throated Dipper *Cinclus cinclus*. *Journal of Ornithology* 154, 91–97.
- Postma, E., Heinrich, F., Koller, U., Sardell, R. J., Reid, J. M., Arcese, P., Keller, L. F., 2011. Disentangling the effect of genes, the environment and chance on sex ratio variation in a wild bird population. *Proceedings of the Royal Society B: Biological Sciences* 278 (1720), 2996–3002.
- R Development Core Team, 2008. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0.
- Robert, C., 2001. The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation. Springer Texts in Statistics. Springer.
- Verner, J., 1965. Selection for sex ratio. *The American Naturalist* 99, 419–421.
- West, S. A., 2009. Sex allocation. Vol. 44. Princeton University Press.
- West, S. A., Herre, E. A., 1998. Stabilizing selection and variance in fig wasp sex ratios. *Evolution* 52, 475–485.
- Witethom, B., Gordh, G., 1994. Development and life table of *Goniozus thailandensis* Gordh & Witethom (Hymenoptera: Bethyridae), a gregarious ectoparasitoid of a phycitine fruit borer (Lepidoptera: Pyralidae). *Journal of the Science Society of Thailand* 20, 101–114.